Gerald F. DeJong
University of Illinois at Urbana–Champaign, Beckman Institute
(217) 333–0491
dejong@cs.uiuc.edu
Machine Learning for Flexible Robotics
N00014–86–K–0309 FINAL REPORT
1 October 1986 – 31 December 1990

SUMMARY OF TECHNICAL PROGRESS

Robotic planning, if it is to be successful in real–world situations, must find some way to side–step the now–well–documented obstacles to classical AI planning. These recent results show that the computational complexity of standard planning is unacceptable even with drastic and untenable simplifying assumptions about the world. The source of complexity in real–world robotic domains includes the problems of data uncertainty, large amounts of data to consider, as well as the problem of tractably producing plans according to the given domain rules. Pretending that these complexities do not exist relegates a computer system to a trivialized micro–world with little hope of applications to the real world.

The research of this grant has been directed towards dealing with the real–world constraints that artificial intelligence robotics systems must address. We have made significant progress on two fronts. The first investigates an integrated approach to planning wherein a classical *a priori* planner is augmented with reactive abilities. Classical planning enables the construction of provably–correct plans, but requires perfect a priori information to do so. In the real world, a perfect characterization of a domain is likely to be intractable, and thus classical planning has been limited to toy domains. Reactivity provides sensitivity to the execution environment and flexibility in planning. However, being essentially a hill–climbing approach to planning, it can neither guarantee eventual nor efficient goal achievement. Ours represents an integration of the two approaches to gain their individual strengths while dealing with their limitations. This integration provides a classical planner with the ability to defer goals which it can prove achievable, and to then address these goals during execution. From reactive planning is gained the ability to utilize runtime information, eliminating the need for perfect *a priori* information. As an augmented classical planner, this integrated approach retains the goal–directedness afforded by a priori planning. Through the achievability constraint, this integrated approach also retains the provably–correct nature of plans constructed in classical planning. We have implemented a prototype system to aid in our investigations. Our progress has been documented in the technical literature.

The underlying unification is performed by a new planning technique we developed called contingent explanation–based learning. As in the second approach, learning is a key ingredient. In contingent EBL, deferred goals are represented using conjectured variables, which act as placeholders for the eventual values of plan parameters whose values are unknown prior to execution. Contingent EBL extends traditional EBL by allowing a planner to distinguish between decisions made prior to execution and decisions made during execution. This is done through the use of conjectured variables, which act as placeholders for deferred planning decisions. Contingent EBL also incorporates completors into general plans. Completors are responsible for the runtime determination of values to replace the conjectured variables.. Only conjectured variables with accompanying achievability proofs are allowed into contingent explanations, and thus the general plans learned in contingent EBL are guaranteed to be completable.

Gerald F. DeJong
University of Illinois at Urbana–Champaign, Beckman Institute
(217) 333–0491
dejong@cs.uiuc.edu
Machine Learning for Flexible Robotics
N00014–86–K–0309 FINAL REPORT
1 October 1986 – 31 December 1990

The second thrust of this grant explores a new approach called *permissive* planning. We have implemented our ideas in the GRASPER system which has capabilities to monitor execution of its plans and to tune its model of the world on failure through use of explicit approximations. We have tested GRASPER extensively on two domains: achieving stable robotic grasping of novel objects, and tray–tilting to orient parts (a domain developed at CMU that we adopted for comparison purposes). Approximations, as employed by GRASPER, involve tunable continuous quantities. These approximations include *controls*, which directly control some parameter in the world, *constraints*, which help to evaluate the choice of one from a set of candidates, and *weights*, which evaluate constraints with respect to each other. The system is being employed in the robotic grasping domain to improve its performance at grasping children's puzzle pieces. Approximations are tuned on failure to yield better future performance. This has the effect, with failures due to uncertainty, of reducing failures by tuning approximate rules so as to be more uncertainty tolerant.

Statement A per telecon
Dr. Andre Van Tilborg ONR/Code 1133
Arlington, VA 22217-5000

NWW 6/26/92

Gerald F. DeJong
University of Illinois at Urbana–Champaign, Beckman Institute
(217) 333–0491
dejong@cs.uiuc.edu
Machine Learning for Flexible Robotics
N00014–86–K–0309 FINAL REPORT
1 October 1986 – 31 December 1990

LIST OF PUBLICATIONS, PRESENTATIONS, AND REPORTS

Papers Published in Refereed Journals

1990

J. W. Shavlik and G. F. DeJong, "Learning in Mathematically Oriented Domains," *Artificial Intelligence*, to appear in 1990.

Technical Reports and Working Papers

1990

A. Blanco, "Towards Intelligent Finite Element Analysis," M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, September 1990. (Also appears as Technical Report UIUCDCS–R–90–1626, Department of Computer Science, University of Illinois at Urbana–Champaign.)

G. F. Dejong, "Explanation–Based Learning with Plausible Inference," Technical Report UIUCDCS–R–90–1577, Department of Computer Science, University of Illinois, Urbana, IL, March 1990.

M. T. Gervasio, "Learning Completable Reactive Plans Through Achievability Proofs," M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, May 1990. (Also appears as Technical Report UIUCDCS–R–90–1605, Department of Computer Science, University of Illinois at Urbana–Champaign.)

1989

S. W. Bennett, "Learning Uncertainty Tolerant Plans Through Approximation in Complex Domains," M.S. Thesis, ECE, University of Illinois, Urbana, Il., January 1989. (Also appears as Technical Report UILU–ENG–89–2204, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign)

S. A. Rajamoney, "Explanation–Based Theory Revision: An Approach to the Problems of Incomplete and Incorrect Theories," Ph.D. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, December 1988.

1988

S. W. Bennett, "An Approach to Learning and Refining Error Tolerant Plans," Working Paper 86, AI Research Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL., October 1988.

S. A. Chien, "Using and Refining Simplifications: Explanation–based Learning of Plans in Intractable Domains," Working Paper 85, AI Research Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL., March 1988.

R. J. Mooney, "A General Explanation–Based Learning Mechanism and its Application to Narrative Understanding," Ph.D. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, January 1988. (Also appears as UILU–ENG–87–2269, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

J. W. Shavlik, "Generalizing the Structure of Explanations in Explanation–Based Learning," Ph.D. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, January 1988. (Also appears as UILU–ENG–87–2276, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

1987

S. A. Chien, "On Resource Constrained Inference," Working Paper 82, AI Research Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL., February 1987.

S. A. Chien, "A Framework for Using Abstraction in Learning," Working Paper 83, AI Research Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL., February 1987.

S. A. Chien, "Simplifications in Temporal Persistence: An Approach to the Intractable Domain Theory Problem in Explanation–Based Learning," M.S. Thesis, Department of Computer Science, University of Illinois, Urbana, IL, August 1987. (Also appears as UILU–ENG–87–2255, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

R. J. Mooney and G. F. DeJong, "Learning Indices for Conceptual Information Retrieval," Technical Report UILU–ENG–87–2230, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign, Urbana, IL, May 1987.

S. A. Rajamoney and G. F. DeJong, "ACTIVE AMBIGUITY REDUCTION: An Experiment Design Approach to Tractable Qualitative Reasoning," Technical Report UILU–ENG–87–2225, April 1987.

J. W. Shavlik and G. F. DeJong, "Acquiring General Iterative Concepts by Reformulating Explanations of Observed Examples," Technical Report UILU–ENG–87–2277, December 1987. (This is an extended version of a chapter appearing in Machine Learning: An Artificial Intelligence Approach, Volume III, R. Michalski and Y. Kodratoff (eds.), Morgan–Kaufmann, 1988.)

1986

S. A. Chien, "A Failure–Driven Approach to Schema Refinement," Working Paper 81, AI Research Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL, 1986.

Invited Presentations

1990

G. F. DeJong and S. W. Bennett, "Learning to Plan in Uncertain and Continuous Domains," *Proceedings of the Office of Naval Research Workshop on Knowledge Acquisition*, Arlington, VA, November 1989.

Gerald F. DeJong
University of Illinois at Urbana–Champaign, Beckman Institute
(217) 333–0491
dejong@cs.uiuc.edu

G. DeJong, "A Machine Learning Approach to Intelligent Adaptive Control," *Proceedings of the 29th Institute of Electrical and Electronics Engineers Conference on Decision and Control*, Honolulu, Hawaii, December 1990.

M. T. Gervasio, "Using Qualitative Reasoning in Proving Achievability," *Proceedings of the 29th Institute of Electrical and Electronics Engineers Conference on Decision and Control*, Honolulu, Hawaii, December 1990.

1989

G. F. DeJong, "Explanation–Based Learning with Plausible Inferencing," *Proceedings of The Fourth European Working Session on Learning*, Montpellier, December 1989, pp. 1–10.

1988

G. F. DeJong, "An Overview of Explanation–Based Learning," *Invited Talk, McDonnell–Douglas Aircraft Company*, St. Louis, MO, July 1988.

1987

G. F. DeJong, "Knowledge–Based Learning for Knowledge–Based Systems," *Keynote Address, Readiness 2000, AOG/AAAIC Conference*, Dayton, OH, October 1987.

S. A. Rajamoney and G. F. DeJong, "An Experiment Design Approach to Ambiguity Reduction," *Invited Talk, Workshop for Qualitative Physics*, Urbana, IL, May 1987.

Contributed Presentations

1991

S. Bennett and G. DeJong, "Comparing Stochastic Planning to the Acquisition of Increasingly Permissive Plans for Complex, Uncertain Domains," *Proceedings of the Eighth International Workshop on Machine Learning*, Evanston, IL, June 1991.

S. A. Chien, M. T. Gervasio and G. F. DeJong, "Becoming Decreasingly Reactive: Learning to Deliberate Minimally," *Proceedings of the Eighth International Workshop on Machine Learning*, Evanston, IL, June 1991.

1990

S. W. Bennett, "Learning Approximation–based Uncertainty–tolerant Plans," *Working Notes of the Symposium on Planning in Uncertain, Unpredictable or Changing Environments*, Palo Alto, CA, March 1990, pp. 16–20.

S. W. Bennett, "Reducing Real–world Failures of Approximate Explanation–based Rules," *Proceedings of the Seventh International Conference on Machine Learning*, Austin, TX, June 1990, pp. 226–234.

S. W. Bennett, "An Overview of Approximation as Employed by the GRASPER System," *Working Notes of the American Association for Artificial Intelligence–90 Workshop on the Automatic Generation of Approximations and Abstractions*, Boston, MA, July 1990, pp. 1–12.

Gerald F. DeJong
University of Illinois at Urbana–Champaign, Beckman Institute
(217) 333–0491
dejong@cs.uiuc.edu

S. Bennett, "Planning to Address Uncertainty: An Incremental Approach Employing Learning Through Experience," *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control*, San Diego, CA, November 1990, pp. 313–324.

G. F. DeJong, "Plausible Inference vs. Abduction," *Working Notes of the Symposium on Automated Abduction*, Stanford, CA, March 1990.

J. M. Gratch and G. F. DeJong, "A Framework for Evaluating Search Control Strategies," *Proceedings of the Workshop on Innovative Approaches to Planning, Scheduling and Control*, San Diego, CA, November 1990, pp. 337–347.

M. T. Gervasio, "Learning General Completable Reactive Plans," *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston, MA, August 1990, pp. 1016–1021.

1989

S. Bennett, "Learning Approximate Plans for Use in the Real World," *The Sixth Internation Workshop on Machine Learning*, Ithaca, NY, 1989, pp. 224–228.

S. A. Chien, "Failure–guided Search in Planning," *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Planning and Search*, Palo Alto, CA, March 1989.

S. A. Chien, "Learning by Analyzing Fortuitous Occurrences," *Proceedings of the 1989 International Machine Learning Workshop*, Ithaca, NY, June 1989, pp. 249–251.

S. A. Chien, "Using and Refining Simplifications: Explanation–based Learning of Plans in Intractable Domains," *Proceedings of The Eleventh International Joint Conference on Artificial Intelligence*, Detroit, MI, August 1989, pp. 590–595.

M. T. Gervasio and G. F. DeJong, "Explanation–Based Learning of Reactive Operators," *Proceedings of the Sixth International Workshop on Machine Learning*, Ithaca, NY, June 1989, pp. 252–254.

1988

S. W. Bennett, "Real World EBL: Learning Error Tolerant Plans in the Robotics Domain," *Proceedings of the 1988 American Association for Artificial Intelligence Spring Symposium Series on Explanation–based Learning*, Stanford, CA, March 1988, pp. 122–126.

S. A. Chien, "A Framework for Explanation–based Refinement," *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Explanation–Based Learning*, Palo Alto, CA, March 1988.

S. A. Chien and G. F. DeJong, "Recognizing Prevention in Plans for Explanation–Based Learning," *Proceedings of the American Association for Artificial Intelligence Workshop on Plan Recognition*, Minneapolis, MN, August 1988.

G. F. DeJong, "Some Thoughts on the Present and Future of Explanation–Based Learning," *1988 European Conference on Artificial Intelligence*, August 1988, pp. 690–697.

B. C. Falkenhainer and S. A. Rajamoney, "The Interdependencies of Theory Formation, Revision, and Experimentation," *Proceedings of the Fifth International Conference on Machine*

Learning, Ann Arbor, MI, June 1988. (Also appears as Technical Report UILU–ENG–87–2224, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

R. J. Mooney, "Generalizing the Order of Operators and its Relation to Generalizing Structure," *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Explanation–Based Learning*, Palo Alto, CA, March 1988.

R. J. Mooney, "Generalizing the Order of Operators in Macro–Operators," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, MI, June 1988, pp. 270–283.

S. A. Rajamoney, "Experimentation–based Theory Revision," *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Explanation–Based Learning*, Palo Alto, CA, March 1988.

S. A. Rajamoney and G. F. DeJong, "Active Explanation Reduction: An Approach to the Multiple Explanations Problem," *Proceedings of the Fifth International Conference on Machine Learning*, Ann Arbor, MI, June 1988.

1987

W. Ahn, R. J. Mooney, W. F. Brewer and G. F. DeJong, "Schema Acquisition from One Example: Psychological Evidence for Explanation–Based Learning," *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA, July 1987, pp. 50–57. (Also appears as Technical Report UILU–ENG–87–2231, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign)

S. W. Bennett, "Approximation in Mathematical Domains," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy, August 1987, pp. 239–241. (Also appears as Technical Report UILU–ENG–87–2238, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

S. A. Chien, "Extending Explanation–Based Learning: Failure–Driven Schema Refinement," *Proceedings of the Third IEEE Conference on Artificial Intelligence Applications*, Orlando, Florida, February 1987. (Also appears as Technical Report UILU–ENG–87–2203, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

S. A. Chien, "Incremental Explanation: An Approach to Reducing Complexity in Learning," *Proceedings of The First Annual Meeting of the Midwest Artificial Intelligence and Cognitive Science Society*, Chicago, IL, April 1987.

R. J. Mooney, "Explanation–Based Learning: A General Learning Mechanism and its Application to Several Complex Domains," *Complex Learning Workshop*, Grange–over–Sands, England, April 1987.

R. J. Mooney, "Integrated Learning of Words and their Underlying Concepts," *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA, July 1987. (Also

appears as Technical Report UILU–ENG–87–2229, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign)

S. A. Rajamoney, "A Model of Experiment Design," *Proceedings of The First Annual Meeting of the Midwest Artificial Intelligence and Cognitive Science Society*, Chicago, IL, April 1987.

S. Rajamoney and G. DeJong, "The Classification, Detection and Handling of Imperfect Theory Problems," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy, August 1987, pp. 205–207. (Also appears as Technical Report UILU–ENG–87–2224, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

1986

R. J. Mooney and S. W. Bennett, "A Domain Independent Explanation–Based Generalizer," *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, August 1986, pp. 551–555. (Also appears as Technical Report UILU–ENG–86–2216, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana–Champaign.)

Books and Book Sections

1989

G. F. DeJong, "The Role of Explanation in Analogy or the Curse of an Alluring Name,"in *Similarity and Analogical Reasoning*, S. Vosniadou and A. Ortony (ed.), Cambridge University, 1989.

1988

G. F. DeJong, "An Introduction to Explanation–Based Learning,"in *Exploring Artificial Intelligence*, H. Shrobe and American Association for Artificial Intelligence (ed.), Morgan Kaufman, San Mateo, CA, 1988.

Gerald F. DeJong
University of Illinois at Urbana–Champaign, Beckman Institute
(217) 333–0491
dejong@cs.uiuc.edu
Machine Learning for Flexible Robotics
N00014–86–K–0309 FINAL REPORT
1 October 1986 – 31 December 1990

DOD INTERACTIONS

DOD interactions have centered around conferences and workshops. Scott Bennett, a graduate student, and the PI attended and presented at the Naval Research Workshop on Knolwedge Acquisition in November, 1989. In August, the PI attended the DOD workshop on Autonomous Vehicles in Crystal City. Several interactions involved personnel from the Naval Research Laboratory and our research group at the AAAI and Machine Learning conference, and the Machine Learning Program Committee meeting.

Gerald F. DeJong
University of Illinois at Urbana–Champaign, Beckman Institute
(217) 333–0491
dejong@cs.uiuc.edu
Machine Learning for Flexible Robotics
N00014–86–K–0309 FINAL REPORT
1 October 1986 – 31 December 1990

DESCRIPTION OF SOFTWARE/HARDWARE PROTOTYPES

Two prototypes have resulted from this research. The first demonstrates a Machine Learning AI system connected to a scara–type RTX robot arm and a simple camera. The system is presented with real–world objects of rather complex shapes. These were not manufactured by the research team but are from a children's plastic jigsaw–like puzzle. The system has no prior knowledge of the piece shapes or information on how to grasp them. After taking a picutre and approximating their shape, the system postulates a grasping strategy and tries it out with the robot manipulator. When the initial strategy fails (because of uncertain or approximated information) the system diagnoses the failure and, using the explanation, propagates the qualitative uncertainty from the real–world to parameters within the grasping strategy. A new strategy is produced that is as uncertainty–tolerant as possible for the diagnosed class of uncertainties. The resulting strategy is available together with sufficient indexing information for future grasping episodes. The prototype system has been run in several data–collection experiments, the empirical findings confirm the theoretical predictions of continuously improved grasping success on novel objects. The second system demonstrates a combined classical/reactive planner of the sort described earlier. It is far less developed but nonetheless demonstrates a proof of concept that *reactive* operators can co–exist within standard operators in a classical planner. An experiment with the prototype demonstrates that, under certain circumstances, such an integrated approach provides significant improvement over classical planning alone, while preserving the domain–independence advantages of classical planning over a purely reactive approach.